

A Multiple Disk, Variable RPM Data Storage System for Reducing Power Consumption

ABSTRACT

5           A data storage system includes a set of disks where each disk is operable in a plurality of discrete angular velocity levels. A disk controller controls the angular velocity of each active disk. The controller replicates a first portion of data on a plurality of the disks stores a second class of data in the set of disks without replication. The disk controller routes data requests to one of the active disks based, at least in part, on the current loading of the active disks to  
10           maintain balanced loading on the active disks. The disk controller alters the angular velocity of at least one of the active disks upon detecting that the latency of one or more of the data requests differs from a specified threshold. In this manner, the disk controller maintains the angular velocity of the active disks at approximately the same minimum angular velocity needed to attain acceptable performance. The disk controller may replicate the first portion of data on each of the  
15           disks in the set of disks. The disk controller may balance the loading on the active disks by routing an incoming request to the active disk with the least loading. The disk controller may maintain each of the active disks at approximately the same angular velocity by preventing the angular velocity of any active disk from differing from the angular velocity of any other active disk by more than one discrete level. The disk controller may recognize two or more levels of  
20           request priorities. In this embodiment, the disk controller routes requests of a first priority to an active disk in a first subset of active disks based, at least in part, on the current loading of the disks in the first subset and route requests of a second priority to an active disk in a second subset of active disks based, at least in part, on the current loading of the disks in the second subset.